arXiv:1701.05324v1 [cs.LO] 19 Jan 2017

A Characterisation of Open Bisimulation using an Intuitionistic Modal Logic

Ki Yung Ahn, Ross Horne and Alwen Tiu
School of Computer Science and Engineering, Nanyang Technological University, Singapore
Email: {yaki,rhorne,atiu}@ntu.edu.sg

***Abstract*—Open bisimulation is the greatest bisimulation congruence for the π-calculus. In open bisimulation, free names in processes are treated as variables that may be instantiated; in contrast to late bisimulation where free names are constants. An established modal logic due to Milner, Parrow, and Walker characterises late bisimulation, that is, two processes satisfy the same set of formulae if and only if they are bisimilar. We propose an intuitionistic variation of this modal logic and prove that it characterises open bisimulation. The soundness proof is mechanised in Abella. The completeness proof provides an algorithm for generating distinguishing formulae, which can be used to explain and certify whenever processes are non-bisimilar.**

I. Introduction

In this work, we consider open bisimulation [1] which ensures processes are equivalent in any context at any point in their execution [2]. There is a practical motivation behind such an equivalence. In modern computing, it is the norm that process can be embedded in unforeseen environments or reused as components in a larger system. When comparing two processes, it is insufficient to know that your methods are sound within the boundaries of your sandbox test environment. Two processes that are observationally equivalent within a sandbox, when unleashed into the world where they can be executed in unforeseen contexts, may exhibit observably distinct behaviours. Such an oversight can introduce bugs difficult to detect, particularly given that the two processes remain "observationally equivalent" within a sandbox.

More subtly still, another norm of modern computing is that processes change context during execution; for example, virtual machines migrate between devices, and replicas replace components at runtime to keep a system live in the face of unavoidable node failures. Two processes may be observationally indistinguishable when executed in any given context. However, if the same two processes execute a few steps and then are migrated to another context, then it is possible that, from that point, the processes can exhibit observably distinct behaviours. If you are testing a program you would never detect this kind of bug even if you test the processes in as many different environments as you can imagine. However, if you then release the processes and someone decides to migrate your process during execution, two "provably equivalent" processes can exhibit observably distinct behaviours.

Process equivalences for π-calculus coarser than open bisimulation are prone to the limitations described above. For instance, late bisimulation [3] is useful only if we can fix the bounds of the entire system. With late bisimulation we can compare two systems, in their entirety, and equate them if at every point in their execution they have the same observable behaviours. The issue is that late bisimulation is not a congruence. Furthermore, even if we take the greatest congruence relation contained in late bisimulation, called *late congruence*, late congruence is no longer a bisimulation hence is not preserved at runtime. These issues are remedied by open bisimulation [1].

The problem that this paper addresses is the nature of a modal logic characterising open bisimulation. A modal logic characterising a bisimulation should have the property that whenever two processes are not bisimilar there should exist a distinguishing formula in the modal logic that holds for one process, but not for the other process. Such distinguishing formulae are useful for explaining why two processes are not bisimilar. Modal logics characterising late bisimulation and other coarser bisimulations were developed early in the literature on the π-calculus, by Milner, Parrow and Walker [4].

In the open bisimulation game between a pair of processes, at every stage, a name substitution is applied before choosing a transition of one of the processes. The other process must apply the same substitution before choosing a matching action. This is illustrated in the following simplified clause for open bisimulation (involving only τ actions, and ignoring name restriction for now): an open bisimulation $\mathcal{R}$ satisfies a clause where, whenever $(P, Q) \in \mathcal{R}$, for all substitutions θ, if $P\theta \xrightarrow{\tau} P'$ there exists Q' such that $Q\theta \xrightarrow{\tau} Q'$ and $(P', Q') \in \mathcal{R}$.

A modal logic for open bisimulation needs to reflect in its semantics the potential substitutions that can be applied to a process. A natural semantics would be a Kripke-like semantics, where worlds are processes and the accessibility relation relates instances of processes. More precisely, consider a relation on processes (worlds) as follows: $P \leq Q$ iff there exists a substitution θ such that $P\theta = Q$. It is then easy to see that the pair $(\mathcal{P}, \leq)$, where $\mathcal{P}$ is the set of processes, forms a Kripke frame that is reflexive and transitive. Consequently, we obtain a semantics for an intuitionistic logic, with implication defined as follows:

$$P \models \varphi_1 \supset \varphi_2 \text{ iff } \forall\theta, \text{ if } P\theta \models \varphi_1\theta, \text{ then } P\theta \models \varphi_2\theta.$$

Intuitionistic negation $\neg\phi$ can then be defined as $\phi \supset \mathtt{ff}$.

To capture transitions of processes, we also require modal operators $\langle\tau\rangle$ and $[\tau]$. The operator $\langle\tau\rangle$ is defined as in the modal logic for late bisimulation, i.e.,

$$P \models \langle\tau\rangle\varphi \text{ iff } \exists P', P \xrightarrow{\tau} P' \text{ and } P' \models \varphi.$$

It is not difficult to check that the law of excluded middle does not hold in this logic, for example, we have:

$$\overline{a}b \parallel c(x) \not\models \langle\tau\rangle\mathtt{tt} \vee \neg\langle\tau\rangle\mathtt{tt}.$$

Neither $\overline{a}b \parallel c(x) \models \langle\tau\rangle\mathtt{tt}$ nor $\overline{a}b \parallel c(x) \models \neg\langle\tau\rangle\mathtt{tt}$ holds. The former holds only if $\overline{a}b \parallel c(x)$ is guaranteed to make a τ transition; but such a transition is only possible assuming $a = c$, hence $\overline{a}b \parallel c(x) \not\models \langle\tau\rangle\mathtt{tt}$. For the latter, by definition of implication, we should consider all reachable worlds, including $(\overline{a}b \parallel c(x))\{^c/_a\}$ which can make a τ transition, but $(\overline{a}b \parallel c(x))\{^c/_a\} \not\models \mathtt{ff}$; thereby $\overline{a}b \parallel c(x) \not\models \neg\langle\tau\rangle\mathtt{tt}$. The failure of the above formulae relies on the fact that we don't know yet whether $a = c$ or $a \neq c$, which amounts to rejecting the law of excluded middle for name equality, something that has been observed in related work on logical encodings of open bisimulation [5].

As a further example, consider the following two processes.

$$\begin{array}{l}\tau.(\overline{a}b \parallel c(x)) \\ \quad + \tau.(\overline{a}b.c(x) + c(x).\overline{a}b) \\ \quad + \tau.(\overline{a}b.a(x) + a(x).\overline{a}b + \tau)\end{array} \qquad \begin{array}{l}\tau.(\overline{a}b.c(x) + c(x).\overline{a}b) \\ \quad + \tau.(\overline{a}b.a(x) + a(x).\overline{a}b + \tau)\end{array}$$

They are not open bisimilar. The process on the right satisfies $[\tau](\langle\tau\rangle\mathtt{tt} \vee \neg\langle\tau\rangle\mathtt{tt})$ but the process on the left does not, since there is a τ-transition to the process $\overline{a}b \parallel c(x)$ that we just agreed above does not satisfy $\langle\tau\rangle\mathtt{tt} \vee \neg\langle\tau\rangle\mathtt{tt}$. In this example, the absence of the law of excluded middle is necessary for existence of distinguishing formaulae for these processes.

The fact that our logic is intuitionistic complicates the construction of distinguishing formulae for non-bisimilar processes. For example, τ and $[a = c]\tau$ are not open bisimilar, so there should be a distinguishing formula to witness this. One such formula is $\langle\tau\rangle\mathtt{tt}$, for which $\tau \models \langle\tau\rangle\mathtt{tt}$ and $[a = c]\tau \not\models \langle\tau\rangle\mathtt{tt}$. This particular construction has a bias towards the process τ. In the classical setting of modal logic for late bisimulation, given such a distinguishing formula, we can dualise it to obtain another distinguishing formula $\neg\langle\tau\rangle\mathtt{tt}$ that has a bias towards $[a = c]\tau$, i.e., $[a = c]\tau \models \neg\langle\tau\rangle\mathtt{tt}$ but $\tau \not\models \neg\langle\tau\rangle\mathtt{tt}$. This dual construction **fails** in the case of our intuitionistic logic characterising open bisimulation. In the intuitionistic setting, we have both $\tau \not\models \neg\langle\tau\rangle\mathtt{tt}$ and $[a = c]\tau \not\models \neg\langle\tau\rangle\mathtt{tt}$. To address this problem, our procedure (c.f. Section IV) simultaneously constructs two distinguishing formulae, that are not necessarily dual to each other.

The precise semantics are presented in the body of this paper. The techniques are clean and modular, so results extend to open bisimulation for more expressive process calculi.

Outline: Section II introduces the semantics of Open Milner–Parrow–Walker logic ($\mathcal{OM}$) and states the soundness and completeness results. Section III describes how the proof assistant Abella [6] was used to mechanically prove soundness and also parts of completeness of the logic with respect to open bisimulation. Section IV presents the proof of the correctness of an algorithm for generating distinguishing formulae, which is used to establish completeness of the logic with respect to open bisimulation.

$$\begin{array}{lll} P ::= & 0 & \text{(deadlock)} \\ & \nu x.P & \text{(nu)} \\ & x(y).P & \text{(input)} \\ & \overline{x}y.P & \text{(output)} \\ & \tau.P & \text{(progress)} \\ & [x = y]P & \text{(match)} \\ & P \parallel P & \text{(par)} \\ & P + P & \text{(choice)} \end{array} \qquad \begin{array}{lll} h ::= & \epsilon & \text{(empty)} \\ & h \cdot x^o & \text{(name)} \\ & h \cdot x^i & \text{(variable)} \end{array}$$

$$\begin{array}{lll} \pi ::= & \tau & \text{(progress)} \\ & \overline{x}z & \text{(free out)} \\ & \overline{x}(z) & \text{(bound out)} \\ & x(z) & \text{(input)} \end{array} \qquad \begin{array}{lll} \phi ::= & \mathtt{tt} & \text{(true)} \\ & \mathtt{ff} & \text{(false)} \\ & \langle\pi\rangle\phi & \text{(dia-action)} \\ & [\pi]\phi & \text{(box-action)} \\ & \langle x = y\rangle\phi & \text{(dia-match)} \\ & [x = y]\phi & \text{(box-match)} \\ & \phi \wedge \phi & \text{(and)} \\ & \phi \vee \phi & \text{(or)} \\ & \phi \supset \phi & \text{(implies)} \end{array}$$

$$\frac{}{x(y).P \xrightarrow{x(y)} P} \qquad \frac{}{\overline{x}y.P \xrightarrow{\overline{x}y} P}$$

$$\frac{P \xrightarrow{\pi} Q}{\nu x.P \xrightarrow{\pi} \nu x.Q}\ x \notin \mathrm{n}(\pi) \qquad \frac{P \xrightarrow{\pi} Q}{P \parallel R \xrightarrow{\pi} Q \parallel R}\ \begin{array}{l}\text{if } \pi = \overline{x}(z) \\ \text{or } \pi = x(z) \\ \text{then } z\#R\end{array}$$

$$\frac{P \xrightarrow{\overline{x}(z)} P' \quad Q \xrightarrow{x(z)} Q'}{P \parallel Q \xrightarrow{\tau} \nu z.(P' \parallel Q')} \qquad \frac{P \xrightarrow{\overline{x}y} P' \quad Q \xrightarrow{x(z)} Q'}{P \parallel Q \xrightarrow{\tau} P' \parallel Q'\{^y/_z\}}$$

$$\frac{P \xrightarrow{\overline{x}z} Q}{\nu z.P \xrightarrow{\overline{x}(z)} Q}\ x \neq z \qquad \frac{P \xrightarrow{\pi} R}{P + Q \xrightarrow{\pi} R} \qquad \frac{P \xrightarrow{\pi} R}{[x = x]P \xrightarrow{\pi} R}$$

where $n(x(y)) = n(\overline{x}(y)) = n(\overline{x}y) = \{x, y\}$ and $n(\tau) = \emptyset$;
and $x\#P$ is such that x is fresh for P
(note, $\nu x.P$ and $z(x).P$ bind x in P).

- $P \models^h \mathtt{tt}$ and $P \not\models^h \mathtt{ff}$ always hold.
- $P \models^h \phi_1 \wedge \phi_2$ iff $P \models^h \phi_1$ and $P \models^h \phi_2$.
- $P \models^h \phi_1 \vee \phi_1$ iff $P \models^h \phi_1$ or $P \models^h \phi_2$.
- $P \models^h \phi_1 \supset \phi_2$ iff
  $\forall\sigma$ respects h, $P\sigma \models^{h\sigma} \phi_1\sigma \implies P\sigma \models^{h\sigma} \phi_2\sigma$.
- $P \models^h \langle x = x\rangle\phi$ iff $P \models^h \phi$.
- $P \models^h \langle\alpha\rangle\phi$ iff $\exists\, Q$, $P \xrightarrow{\alpha} Q$ and $Q \models^h \phi$.
- $*$ $P \models^h \langle\overline{a}(x)\rangle\phi$ iff $\exists\, Q$, $P \xrightarrow{\overline{a}(x)} Q$ and $Q \models^{h\cdot x^o} \phi$.
- $*$ $P \models^h \langle a(x)\rangle\phi$ iff $\exists\, Q$, $P \xrightarrow{a(x)} Q$ and $Q \models^{h\cdot x^i} \phi$.
- $P \models^h [x = y]\phi$ iff
  $\forall\sigma$ respects h, $x\sigma = y\sigma \implies P\sigma \models^{h\sigma} \phi\sigma$.
- $P \models^h [\alpha]\phi$ iff
  $\forall\sigma$ respects $h, \forall Q$, $P\sigma \xrightarrow{\alpha} Q \implies Q \models^{h\sigma} \phi\sigma$.
- $*$ $P \models^h [\overline{a}(x)]\phi$ iff
  $\forall\sigma$ respects h, $\forall Q$, $P\sigma \xrightarrow{\overline{a}(x)} Q \implies Q \models^{h\sigma\cdot x^o} \phi\sigma$.
- $*$ $P \models^h [a(x)]\phi$ iff
  $\forall\sigma$ respects h, $\forall Q$, $P\sigma \xrightarrow{a(x)} Q \implies Q \models^{h\sigma\cdot x^i} \phi\sigma$.

where α is action τ or $\overline{a}b$;
and x is fresh for P, h, and σ in the asterisk($*$)-bullet items.

Fig. 1. Syntax and semantics of the π-calculus and the modal logic $\mathcal{OM}$. Implication ($\supset$) is only used in the motivating discussion in Section I, which explains the intuitionistic nature of the logic; but is not necessary for establishing results, hence is not considered as a part of $\mathcal{OM}$ from Section II.

II. Open Milner–Parrow–Walker Logic ($\mathcal{OM}$)

We recall the syntax and labelled transition semantics for the finite π-calculus (Fig. 1). All features are standard: the deadlocked process that can do nothing, the ν quantifiers that binds private names, the output prefix that outputs a name on a channel, the input prefix that binds the name received on a channel, the silent progress action τ, the name match guard, parallel composition and non-deterministic choice.

There are four types of action ranged over by π, where a *free output* sends a free name, whereas a *bound output* extrudes a ν-bounded private name. Stylistically, the semantics is the late labelled transition system for the π-calculus, where the name on the input channel is a symbolic place holder for a name that is not chosen until after an input transition.

Histories are used to define both the intuitionistic modal logic and open bisimulation. Histories represent what is known about free variables due to how they have been communicated previously to the environment. There are two types of event to record in a history: The output of a fresh private name, using action $\overline{a}(x)$, which is denoted x^o; and a (symbolic) input, using action $a(z)$, which is denoted z^i. The only thing that matters about the order of events in the history is the alternation between the bound outputs and symbolic inputs, since an input variable can only be instantiated with private names that were output earlier in the history. E.g. for history $x^o \cdot z^i$, input variable z may be instantiated with private name x; in contrast, for history $z^i \cdot x^o$, input variable z may not be instantiated with private name x. This is reflected by the constraints on substitutions in the following definition.

Definition 2.1 (σ respects h): A substitution σ respects a history h according to the following inductive definition.

$$\frac{}{\sigma \text{ respects } \epsilon} \qquad \frac{\sigma \text{ respects } h}{\sigma \text{ respects } h \cdot x^i} \qquad \frac{x \notin \mathrm{dom}(\sigma) \cup \mathrm{fv}(h\sigma) \quad \sigma \text{ respects } h}{\sigma \text{ respects } h \cdot x^o}$$

Note that the above inductive definition fulfils the role of sets of inequality constraints called *distinctions* in the original work on open bisimulation [1]. The definition above also captures the alternations between nominal and universal quantifiers in embeddings of open bisimulation in intuitionistic nominal logic LINC [5], [7]. Although distinctions are more general than histories, it is shown in [5] that given a history h and its corresponding distinction D, the corresponding definitions of open bisimilarity coincide.

a) The Modal Logic $\mathcal{OM}$: The semantics of the modal logic $\mathcal{OM}$ is defined in terms of the late labelled transitions system (Fig. 1) and history respecting substitutions (Definition 2.1). Intuitively, each judgement must hold for all possible respectful substitutions, which explains the asymmetry between the box and diamond modalities. For a diamond modality the action inside the modality must be possible, regardless of the substitution. Hence, since applying a respectful substitution cannot prevent an action occurring, it is sufficient to consider the identity substitution. In contrast, for the box modality there may be zero or more distinct substitutions enabling the action and we should consider all of them, explaining the quantification over all substitutions.

Definition 2.2 ($P \models^h \phi$): A process P satisfies a formula ϕ with history h according to the inductive definition in Fig. 1. Satisfaction, written $P \models \phi$, is satisfaction with a history of inputs $x_0^i \cdot \ldots \cdot x_n^i$, where $\mathrm{fv}(P) \cup \mathrm{fv}(Q) \subseteq \{x_0, \ldots, x_n\}$.

A. Open Bisimulation, Soundness and Completeness

We recall the definition of open bisimulation. Open bisimulation is a greatest fixed point of symmetric relations closed under all respectful substitutions and labelled transitions actions at every step. Notice that a symbolic input or output of a fresh private name updates the history.

Definition 2.3 ($P \sim^h Q$): Open bisimilarity with history h is the greatest symmetric relation such that for all substitutions σ such that σ respects h, the following hold, where α is a τ or $\overline{a}b$ action and x is fresh for $P\sigma$, $Q\sigma$ and $h\sigma$:

- $P\sigma \xrightarrow{\alpha} P' \implies \exists\, Q',\ Q\sigma \xrightarrow{\alpha} Q'$ and $P' \sim^{h\sigma} Q'$.
- $P\sigma \xrightarrow{\overline{a}(x)} P' \implies \exists\, Q',\ Q\sigma \xrightarrow{\overline{a}(x)} Q'$ and $P' \sim^{h\sigma \cdot x^o} Q'$.
- $P\sigma \xrightarrow{a(x)} P' \implies \exists\, Q',\ Q\sigma \xrightarrow{a(x)} Q'$ and $P' \sim^{h\sigma \cdot x^i} Q'$.

Open bisimilarity, written $P \sim Q$, is open bisimilarity with a history $x_0^i \cdot \ldots \cdot x_n^i$, where $\mathrm{fv}(P) \cup \mathrm{fv}(Q) \subseteq \{x_0, \ldots, x_n\}$.

The condition on free variables for $\sim$, where the history is implicit, indicates that the values of all free variables cannot be influenced by private names output in the future.

b) Open bisimulation examples: We recall properties of open bisimulation and classic examples [1]. Open bisimulation has desirable properties that address limitations of late bisimulation first proposed for the π-calculus [4].

Open bisimulation is a congruence relation, i.e. a relation preserved in all contexts. In contrast, late bisimulation is not preserved under input prefix. For example, 0 and $[x = y]\overline{c}d$ are distinguished by the context $z(x).\{\ \cdot\ \} \parallel \overline{z}y$. Open bisimulation correctly distinguishes these processes, whereas late bisimulation identifies these processes.

The above examples could however be distinguished by late bisimulation congruence, which is the least congruence relation containing late bisimulation (checked by considering late bisimulation under any substitution). For a more subtle example, consider the process $P \triangleq \tau.[x = y]\tau + \tau.\tau + \tau$ which is not open bisimilar to $Q \triangleq \tau.\tau + \tau$. To see that P and Q are not open bisimilar observe that transition $P \xrightarrow{\tau} [x = y]\tau$ can only be matched either by $Q \xrightarrow{\tau} 0$ or $Q \xrightarrow{\tau} \tau$. In the former case, by applying a substitution $[x = y]\tau\{y/x\} \xrightarrow{\tau} 0$, but $0\{y/x\}$ can make no τ-transition. In the latter case, $\tau \xrightarrow{\tau} 0$ but $[x = y]\tau$ is not guaranteed to make a τ-transition. In terms of contexts, there is no context distinguishing P from Q; however, after one internal communication, we reach two pairs of states both of which have a distinguishing context.

c) Soundness and completeness results: The main result of this paper is that, for finite π-calculus processes open bisimulation ($\sim$) coincides with the relation between processes with no distinguishing formula ($\overset{\models}{\sim}$).

Definition 2.4 ($P \overset{\models}{\sim} Q$): For all ϕ, $P \models \phi$ iff $Q \models \phi$.

Theorem 2.5 (Soundness): For π-calculus processes (including replication), $P \sim Q$ implies $P \overset{\models}{\sim} Q$.

Theorem 2.6 (Completeness): For finite π-calculus processes, $P \stackrel{\text{ltsf}}{\sim} Q$ implies $P \sim Q$.

The proof of soundness, which has been mechanically checked in the proof assistant Abella [6], is postponed until Section III. The proof of completeness is explained in detail in Section IV. Notice that soundness extends to infinite π-calculus processes with replication, but completeness does not in general.

Firstly, we provide examples demonstrating the implications of Theorems 2.5 and 2.6. Due to soundness, if two processes are bisimilar, we cannot find a distinguishing formula that holds for one process but does not hold for the other process. Due to completeness, if it is impossible to prove that two process are open bisimilar, then we can construct a distinguishing formula that holds for one process but does not hold for the other process. Thus Theorems 2.5 and 2.6 guarantee that an $\mathcal{OM}$ formulae can be used to characterise non-bisimilarity.

d) Example processes distinguishable by postconditions: All modalities are essential for the soundness and completeness of $\mathcal{OM}$. Perhaps the least obvious modality is $\langle x = y\rangle$. When prefixed with a box modality it indicates a postcondition that always holds after an action. To see, this consider the process $[x = y]\tau$. The judgement $[x = y]\tau \models [\tau]\langle x = y\rangle\mathtt{tt}$ holds since for any θ such that $([x = y]\tau)\theta \xrightarrow{\tau} 0$ it must be the case that $x\theta = y\theta$. Hence, by definition of diamond, $0 \models \langle x\theta = y\theta\rangle\mathtt{tt}$ iff $0 \models \mathtt{tt}$. In contrast, $\tau \not\models [\tau]\langle x = y\rangle\mathtt{tt}$ since, taking the identity substitution in the definition of box, $\tau \xrightarrow{\tau} 0$, but $0 \models \langle x = y\rangle\mathtt{tt}$ cannot be proven in general. The formula $[\tau]\langle x = y\rangle\mathtt{tt}$ is therefore a distinguishing formula satisfy $[x = y]\tau$ but not τ.

The use of $\langle x = y\rangle$ as a postcondition contrasts to the use of $[x = y]$ as a precondition. Consider the same process as above with the formula $[x = y]\langle\tau\rangle\mathtt{tt}$. Observe that, for substitutions θ such that $x\theta = y\theta$, $([x = y]\tau)\theta \xrightarrow{\tau} 0$ and $0 \models \mathtt{tt}$ holds, hence $([x = y]\tau)\theta \models \langle\tau\rangle\mathtt{tt}$; and thereby the judgement $[x = y]\tau \models [x = y]\langle\tau\rangle\mathtt{tt}$ holds. In contrast, $0 \not\models [x = y]\langle\tau\rangle\mathtt{tt}$.

We will return to the above two formulae shortly, as they are critical for the algorithm for generating distinguishing formulae.

B. Sketch of Algorithm for Distinguishing Formulae

The completeness proof in Section IV relies on an algorithm for generating distinguishing formulae for non-bisimilar processes. Here, we provide a sketch of the algorithm executed on key examples.

e) Example requiring intuitionistic assumptions: The algorithm proceeds over the structure of a tree of moves that show two processes are non-bisimilar. In a base case, we have a pair of processes where one can perform an action under a substitution that the other process cannot perform under the same substitution. We revisit two examples of the base cases:

$[x = y]\tau \nsim 0$: The left leads by $[x = y]\tau\{^{y}/_{x}\} \xrightarrow{\tau} 0$, but 0 cannot perform a τ-action under any substitution. This a base case, hence $[x = y]\tau \models [x = y]\langle\tau\rangle\mathtt{tt}$ and $0 \models [\tau]\mathtt{ff}$ are distinguishing formulae.

$[x = y]\tau \nsim \tau$: The right leads by $\tau \xrightarrow{\tau} 0$, but $[x = y]\tau\theta$ can perform a τ-action only when $x\theta = y\theta$. This is a base case, hence $[x = y]\tau \models [\tau]\langle x = y\rangle\mathtt{tt}$ and $\tau \models \langle\tau\rangle\mathtt{tt}$ are distinguishing formulae.

In an inductive case, the two processes cannot be distinguished by an immediate action. However, under some substitutions, one process can perform an action to a state, say P', that, under the same substitution the other process can only perform the action to reach states, say Q'_i, that are non-bisimilar to P'. This allows a distinguishing formula to be inductively constructed from the distinguishing formulae for P' paired with each Q'_i.

For example, consider how to algorithmically discover distinguishing formulae for the following processes P and Q.

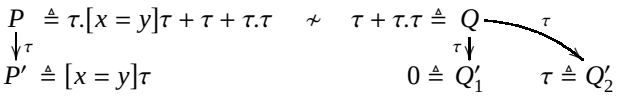

The first step in the strategy for non-bisimilarity is to show that P can perform an action to a state that is not bisimilar to any state reachable by a τ-action for the other process. One possibility is the transition to P' as illustrated above. In reply, Q can make two corresponding transitions to Q'_1 and Q'_2.

Inductively, we require that $P' \nsim Q'_1$ and $P' \nsim Q'_2$. Both are instances of the base case discussed above where we discovered distinguishing formulae for each of them.

This enables us to construct distinguishing formulae for the inductive case. The distinguishing formula satisfying P is a diamond followed by the conjunction of the left distinguishing formulae that satisfy P' in the base cases:

$\tau.[x = y]\tau + \tau + \tau.\tau \models \langle\tau\rangle([\tau]\langle x = y\rangle\mathtt{tt} \wedge [x = y]\langle\tau\rangle\mathtt{tt})$. The distinguishing formula satisfying Q is a box followed by the disjunction of the right distinguishing formulae from the base cases: $\tau + \tau.\tau \models [\tau](\langle\tau\rangle\mathtt{tt} \vee [\tau]\mathtt{ff})$.

To confirm that they are indeed distinguishing formulae for P and Q, swap the processes and formulae above to observe that each formula fails to satisfy the other process. To be precise, assume for contradiction that $\tau + \tau.\tau \models \langle\tau\rangle([\tau]\langle x = y\rangle\mathtt{tt} \wedge [x = y]\langle\tau\rangle\mathtt{tt})$ holds. By definition of $\langle\tau\rangle$, this holds iff either $0 \models [\tau]\langle x = y\rangle\mathtt{tt} \wedge [x = y]\langle\tau\rangle\mathtt{tt}$ or $\tau \models [\tau]\langle x = y\rangle\mathtt{tt} \wedge [x = y]\langle\tau\rangle\mathtt{tt}$ holds. Now observe that both $0 \models [\tau]\langle x = y\rangle\mathtt{tt}$ and $0 \models [x = y]\langle\tau\rangle\mathtt{tt}$ hold iff $x \neq y$. In addition, observe that both $\tau \models [\tau]\langle x = y\rangle\mathtt{tt}$ and $\tau \models [x = y]\langle\tau\rangle\mathtt{tt}$ hold iff $x = y$. Thereby, the assumption above holds iff $\forall x, y.\, (x = y \vee x \neq y)$, which is the law of excluded middle. Fortunately, intuitionistic logics do not assume the law of excluded middle leading to a contradiction. Thereby $Q \not\models \langle\tau\rangle([\tau]\langle x = y\rangle\mathtt{tt} \wedge [x = y]\langle\tau\rangle\mathtt{tt})$, as required.

With the law of excluded middle, the formula for Q also holds for P; and vice versa. Indeed there would be no distinguishing formulae for these processes; and hence in a classical framework the modal logic would be incomplete.

f) Examples involving private names: The alternation between inputs and outputs in the history affects what counts as a respecting substitution. Intuitively, respecting substitutions ensure that a private name can never be input earlier than it

was output. Consider the following processes P and Q.

$$P \triangleq \nu x.\overline{a}x.a(y).\tau \quad \not\sim \quad \nu x.\overline{a}x.a(y).[x=y]\tau \triangleq Q$$

These processes are not open bisimilar because P can perform the following three transition steps:

$$\nu x.\overline{a}x.a(y).\tau \xrightarrow{\overline{a}(x)} a(y).\tau \xrightarrow{a(y)} \tau \xrightarrow{\tau} 0.$$

However, Q can only match the first two steps. At the third step, a base case of the algorithm $\tau \not\sim^{a^i x^o y^i} [x=y]\tau$ applies. In this case, any substitution θ respecting $a^i x^o y^i$ where $[x=y]\tau\theta \xrightarrow{\tau} 0$ is such that $y\theta = x$, $x \notin \mathrm{dom}(\theta)$ and $a\theta \neq x$, which is satisfiable. Thus $[x=y]\tau \models^{a^i x^o y^i} [\tau]\langle x=y\rangle \mathtt{tt}$ and $\tau \models^{a^i x^o y^i} \langle\tau\rangle\mathtt{tt}$. By applying inductive cases, we obtain the following distinguishing formulae:

$$\nu x.\overline{a}x.a(y).\tau \models \langle\overline{a}(x)\rangle\langle a(y)\rangle\langle\tau\rangle\mathtt{tt}$$
$$\nu x.\overline{a}x.a(y).[x=y]\tau \models [\overline{a}(x)][a(y)][\tau]\langle x=y\rangle\mathtt{tt}$$

g) Example with private names but not distinguishable: In contrast to the previous example, consider the following processes where a fresh name is output and compared to a name already known: $\nu x.\overline{a}x \sim \nu x.\overline{a}x.[x=a]\tau$. These processes are open bisimilar, hence by Theorem 2.5 there is no distinguishing formula. The existence of a distinguishing formula of the form $\langle\overline{a}(x)\rangle[x=a]\langle\tau\rangle\mathtt{tt}$ is *prevented* by the history. Both $\nu x.\overline{a}x.[x=a]\tau \models \langle\overline{a}(x)\rangle[x=a]\langle\tau\rangle\mathtt{tt}$ and $\nu x.\overline{a}x \models \langle\overline{a}(x)\rangle[x=a]\langle\tau\rangle\mathtt{tt}$ hold. The latter holds since $\nu x.\overline{a}x \models^{a^i} \langle\overline{a}(x)\rangle[x=a]\langle\tau\rangle\mathtt{tt}$ holds if and only if $\nu x.\overline{a}x \xrightarrow{\overline{a}(x)} 0$ and $0 \models^{a^i x^o} [x=a]\langle\tau\rangle\mathtt{tt}$. By definition of $[x=a]$, this holds if only if for all θ respecting $a^i x^o$ and such that $x\theta = a\theta$, $0 \models^{a^i x^o} \langle\tau\rangle\mathtt{tt}$. Clearly 0 cannot perform a τ-transition, hence $0 \models^{a^i x^o} \langle\tau\rangle\mathtt{tt}$ does not hold. However, fortunately, there is no substitution θ respecting $a^i x^o$ such that $x\theta = a\theta$. By the definition of respecting substitution, θ must satisfy $x \notin \mathrm{dom}(\theta)$ and $x \neq a\theta$, contradicting constraint $x\theta = a\theta$. Thereby $0 \models^{a^i x^o} [x=a]\langle\tau\rangle\mathtt{tt}$ holds vacuously; hence $\nu x.\overline{a}x \models^{a^i} \langle\overline{a}(x)\rangle[x=a]\langle\tau\rangle\mathtt{tt}$ holds as required.

III. Mechanized Proofs in Abella

We developed two versions of the proofs: first with no free variables and second with two free variables. The first helps to grasp the overall structure of our proofs and the second demonstrates that our proofs will generalize for any number of free variables. In general, there are countably many instances of the following theorems to consider, where $\overline{x}$ is the list of n free variables in processes P and Q:

Theorem $\texttt{bisim_sat}_n$: $\forall P\ Q, [\![P(\overline{x}) \sim Q(\overline{x})]\!] \rightarrow [\![P(\overline{x}) \overset{\models}{\sim} Q(\overline{x})]\!]$.

Theorem $\texttt{sat_bisim}_n$: $\forall P\ Q, [\![P(\overline{x}) \overset{\models}{\sim} Q(\overline{x})]\!] \rightarrow [\![P(\overline{x}) \sim Q(\overline{x})]\!]$.

These are the soundness and completeness theorems (Theorems 2.5 and 2.6) in Section II transcribed in Abella. We can develop fully mechanized proofs for $\texttt{bisim_sat}_n$ (Section III-A) and partially mechanized proofs for $\texttt{sat_bisim}_n$ (Section III-B). The missing part of the proof for $\texttt{sat_bisim}_n$ corresponds to Proposition 4.8 in Section IV.

We use the notation $[\![\cdot]\!]$ for the encoding of the concepts defined in Sect. II into Abella's logic. More specifically,

```
sig finite-pic.

kind n  type.

kind p  type.
type null          p.
type taup          p → p.
type plus, par     p → p → p.
type match, out    n → n → p → p.
type in            n → (n → p) → p.
type nu            (n → p) → p.

kind a  type.
type tau           a.
type up, dn        n → n → a.

type one      p →      a →      p  → o.
type oneb     p → (n → a) → (n → p) → o.


module finite-pic.

oneb (in X M) (dn X) M.
one (out X Y P) (up X Y) P.
one  (taup P) tau P.
one  (match X X P) A Q :- one  P A Q.
oneb (match X X P) A M :- oneb P A M.
one  (plus P Q) A R :- one  P A R.
one  (plus P Q) A R :- one  Q A R.
oneb (plus P Q) A M :- oneb P A M.
oneb (plus P Q) A M :- oneb Q A M.
one  (par P Q) A (par P1 Q) :- one P A P1.
one  (par P Q) A (par P Q1) :- one Q A Q1.
oneb (par P Q) A (x\ par (M x) Q)
                              :- oneb P A M.
oneb (par P Q) A (x\ par P (N x))
                              :- oneb Q A N.
one  (nu x\ P x) A (nu x\ Q x)
    :- pi x\ one  (P x) A (Q x).
oneb (nu x\ P x) A (y\ nu x\ Q x y)
    :- pi x\ oneb (P x) A (y\ Q x y).
oneb (nu x\M x) (up X) N
    :- pi y\ one (M y) (up X y) (N y).
one (par P Q) tau (nu y\ par (M y) (N y))
    :- oneb P (dn X) M , oneb Q (up X) N.
one (par P Q) tau (nu y\ par (M y) (N y))
    :- oneb P (up X) M , oneb Q (dn X) N.
one (par P Q) tau (par (M Y) T)
    :- oneb P (dn X) M, one Q (up X Y) T.
one (par P Q) tau (par R (M Y))
    :- oneb Q (dn X) M, one P (up X Y) R.
```

Fig. 2. λProlog specification of the finite π-calculus operational semantics. (Adopted from one of the examples distributed with Abella.)

```
Specification "finite-pic".

CoDefine bisim : p → p → prop
by bisim P Q := (∀ A P₁, {one  P  A     P₁} → ∃ Q₁, {one  Q  A     Q₁} ∧       bisim P₁ Q₁)
              ∧ (∀ X M, {oneb P (dn X) M} → ∃ N, {oneb Q (dn X) N} ∧ ∀ w, bisim (M w) (N w))
              ∧ (∀ X M, {oneb P (up X) M} → ∃ N, {oneb Q (up X) N} ∧ ∇ w, bisim (M w) (N w))
              ∧ (∀ A Q₁, {one  Q  A     Q₁} → ∃ P₁, {one  P  A     P₁} ∧       bisim Q₁ P₁)
              ∧ (∀ X N, {oneb Q (dn X) N} → ∃ M, {oneb P (dn X) M} ∧ ∀ w, bisim (N w) (M w))
              ∧ (∀ X N, {oneb Q (up X) N} → ∃ M, {oneb P (up X) M} ∧ ∇ w, bisim (N w) (M w)).

Kind o′  type.
Type tt, ff          o′.
Type ∨, ∧            o′ → o′ → o′.
Type ⊞, ⊹            n → n → o′ → o′.
Type □, ◇            a → o′ → o′.
Type □↑, ◇↑, □↓, ◇↓  n → (n → o′) → o′.

Define sat : p → o′ → prop
by sat P tt
 ; sat P (∨ A B) := sat P A ∧ sat P B
 ; sat P (∧ A B) := sat P A ∨ sat P B
 ; sat P (⊞ X Y A) := X = Y → sat P A
 ; sat P (⊹ X Y A) := X = Y ∧ sat P A
 ; sat P (□ X A) := ∀ P₁, {one P X P₁} → sat P₁ A
 ; sat P (◇ X A) := ∃ P₁, {one P X P₁} ∧ sat P₁ A
 ; sat P (□↑ X A) := ∀ Q, {oneb P (up X) Q} →  ∇ w, sat (Q w) (A w)
 ; sat P (◇↑ X A) := ∃ Q, {oneb P (up X) Q} ∧  ∇ w, sat (Q w) (A w)
 ; sat P (□↓ X A) := ∀ Q, {oneb P (dn X) Q} →  ∀ w, sat (Q w) (A w)
 ; sat P (◇↓ X A) := ∃ Q, {oneb P (dn X) Q} ∧  ∀ w, sat (Q w) (A w).
```

Fig. 3. A coinductive definition of the bisimulation and an inductive definition of a modal logic for the finite calculus in Abella. (Adopted from one of the examples distributed with Bedwyr [7], which implements [5], and modified for two-level logic style in Abella.)

```
Theorem bisim_sat0 : ∀ P Q F, bisim P Q → ((sat P F → sat Q F) ∧ (sat Q F → sat P F)).
Theorem sat_bisim0 : ∀ P Q, (∀ F, (sat P F → sat Q F) ∧ (sat Q F → sat P F)) → bisim P Q.

Theorem sateq_one_exists_L0 : ∀ P Q A P₁, (∀ F, (sat P F → sat Q F) ∧ (sat Q F → sat P F))
  → {one P A P₁} → ∃ Q₁, {one Q A Q₁}.
Theorem sateq_one_L0 :    ∀ P Q A P₁, (∀ F, (sat P F → sat Q F) ∧ (sat Q F → sat P F))
  → {one P A P₁} → ∃ Q₁, {one Q A Q₁} ∧ (∀ F, (sat P₁ F → sat Q₁ F) ∧ (sat Q₁ F → sat P₁ F)).

Theorem bisim_sat2 : ∀ P Q (F : n → n → o′),
                        (∀ X Y, bisim (P X Y) (Q X Y)) →
                        ( (∀ X Y, sat (P X Y) (F X Y) →  ∀ X Y, sat (Q X Y) (F X Y))
                        ∧ (∀ X Y, sat (Q X Y) (F X Y) →  ∀ X Y, sat (P X Y) (F X Y)) ).
Theorem sat_bisim2 : ∀ (X : n) (Y : n) P Q,
                     (∀ F, (∀ X Y, sat (P X Y) (F X Y) →  ∀ X Y, sat (Q X Y) (F X Y))
                         ∧ (∀ X Y, sat (Q X Y) (F X Y) →  ∀ X Y, sat (P X Y) (F X Y)) ) →
                     bisim (P X Y) (Q X Y).

Theorem sateq_one_exists_L2 : ∀ (X : n) (Y : n) P Q A P₁,
                     ( ∀ F, (∀ X Y, sat (P X Y) (F X Y) →  ∀ X Y, sat (Q X Y) (F X Y))
                          ∧ (∀ X Y, sat (Q X Y) (F X Y) →  ∀ X Y, sat (P X Y) (F X Y)) )
                   → {one (P X Y) (A X Y) (P₁ X Y)} → ∃ Q₁, {one (Q X Y) (A X Y) (Q₁ X Y)}.
Theorem sateq_one_L2 : ∀ (X : n) (Y : n) P Q A P₁,
                     ( ∀ F, (∀ X Y, sat (P X Y) (F X Y) →  ∀ X Y, sat (Q X Y) (F X Y))
                          ∧ (∀ X Y, sat (Q X Y) (F X Y) →  ∀ X Y, sat (P X Y) (F X Y)) )
                   → {one (P X Y) (A X Y) (P₁ X Y)} →
                     ( ∃ Q₁, {one (Q X Y) (A X Y) (Q₁ X Y)} ∧
                       ∀ F, (∀ X Y, sat (P₁ X Y) (F X Y) →  ∀ X Y, sat (Q₁ X Y) (F X Y))
                          ∧ (∀ X Y, sat (Q₁ X Y) (F X Y) →  ∀ X Y, sat (P₁ X Y) (F X Y)) ).
```

Fig. 4. Two versions of main theorems and lemmas stated in Abella: one for closed processes and the other for processes with two free-variables.

$$[\![P(\bar{x}) \sim Q(\bar{x})]\!] \triangleq \forall \bar{x},\ \texttt{bisim}\ ([\![P]\!]\bar{x})\ ([\![Q]\!]\bar{x})$$
$$[\![P(\bar{x}) \overset{\models}{\sim} Q(\bar{x})]\!] \triangleq \forall F,\ [\![P(\bar{x}) \models F(\bar{x}) \iff Q(\bar{x}) \models F(\bar{x})]\!]$$
$$\triangleq \forall F,\ [\![P(\bar{x}) \models F(\bar{x})]\!] \to [\![Q(\bar{x}) \models F(\bar{x})]\!]$$
$$\wedge\ [\![Q(\bar{x}) \models F(\bar{x})]\!] \to [\![P(\bar{x}) \models F(\bar{x})]\!]$$
$$[\![P(\bar{x}) \models F(\bar{x})]\!] \triangleq \forall \bar{x},\ \texttt{sat}\ ([\![P]\!]\bar{x})\ ([\![F]\!]\bar{x})$$

where $\bar{x}$ are the free variables of P, Q, and F and the definitions for `bisim` and `sat` are provided in Fig. 3. The free variables $\bar{x}$ are encoded as universally quantified variables and applied to as arguments for higher-order syntax encodings of P, Q, and F. Both the bisimulation and satisfaction defined in Aballa's reasoning logic refer to the λProlog [8] specification of the finite π-calculus (Fig. 2), which are surrounded by curly braces (e.g., `{one P A P1}`). Abella supports two-level logic approach [9] of having a computational specification in λProlog as an object logic and using Abella's reasoning logic to describe and reason about the properties of the specification.

A. Bisimulation implies Logical Equivalence

We prove `bisim_sat0` by splitting it into two lemmas:

```
bisim_sat0_L : ∀ P Q F, [[P ~ Q]] → [[P ⊨ F]] → [[Q ⊨ F]]
bisim_sat0_R : ∀ P Q F, [[P ~ Q]] → [[Q ⊨ F]] → [[P ⊨ F]]
```

Once we prove one of them, the other is proved simply by instantiating P with Q and Q with P. We prove `bisim_sat0_L` by induction on the second argument $[\![P \models F]\!]$.

Then, it is easy to prove $\texttt{bisim_sat}_n$ for any n because the proofs for the lemmas $\texttt{bisim_sat}_n\texttt{_L}$ and `_R` are straightforward corollaries of the lemmas `bisim_sat0_L` and `_R`. For instance, when $n = 2$, the proof script for `bisim_sat2_L` is as simple as follows:

```
Theorem bisim_sat2_L : ∀ P Q (F : n → n → o'),
    (∀ X Y, bisim (P X Y) (Q X Y))
  → (∀ X Y, sat (P X Y) (F X Y))
  → (∀ X Y, sat (Q X Y) (F X Y)).
  intros.
  assert bisim (P X Y) (Q X Y). backchain H1.
  assert sat (P X Y) (F X Y).  backchain H2.
  backchain bisim_sat0_L.
```

The proof script for $\texttt{bisim_sat}_n\texttt{_L}$ would be structurally identical to above, only differing in the number of free variables (i.e., $\texttt{X}_1 \cdots \texttt{X}_n$ instead of just two variables `X Y`). Therefore, the proof of $(\sim \subseteq \overset{\models}{\sim})$ can be fully mechanized in Abella for any given instance of n.

B. Logical Equivalence implies Bisimulation

We can prove `sat_bisim0` assuming six lemmas, which state a closedness property of logical equivalence with regards to a commonly labeled transition step between two processes. Each of these six closedness lemmas discharges one of the six cases in the coinductive definition of `bisim` (see lines 4-9 in Fig. 3). Here is one of the six lemmas:

```
Theorem sateq_one_L0 : ∀ P Q A P1, [[P ≲ Q]]
    → [[P --A--> P1]] → ∃Q1, [[Q --A--> Q1]] ∧ [[P1 ≲ Q1]].
```

Another lemma `sateq_one_R0` is a symmetric counterpart of above, where Q leads the transition rather than P. This pair of lemmas, `sateq_one_L0` and `_R0`, states that for two logically equivalent processes, if one of them has a transition step, then the other also has a transition step with the same label such that the resulting pair of processes at next step are logically equivalent as well. Additionally, there are two more pairs of closedness lemmas: `sateq_oneb_dn_L0` and `_R0` for bounded inputs and `sateq_oneb_up_L0` and `_R0` for bounded outputs.

It is difficult to directly prove the six closedness lemmas. So, we first prove six existence lemmas, which have weaker conclusions than the closedness lemmas. Here is an existence lemma weakened from `sateq_one_L0`:

```
Theorem sateq_one _exists _L0 : ∀ P Q A P1,
  [[P ≲ Q]] → [[P --A--> P1]] → ∃Q1, [[Q --A--> Q1]].
```

We prove this by case analysis over $[\![P \xrightarrow{A} P_1]\!]$. For each case, we show $[\![Q \models \langle A \rangle \mathtt{tt}]\!]$ from $[\![P \models \langle A \rangle \mathtt{tt}]\!]$, which is obvious from $[\![P \xrightarrow{A} P_1]\!]$. From $[\![Q \models \langle A \rangle \mathtt{tt}]\!]$, $\exists Q_1, [\![Q \xrightarrow{A} Q_1]\!]$ is immediate. The proofs for the other five lemmas are similar.

The only gap we need to fill is from the existence lemmas to their corresponding closedness lemmas. This gap can be filled by two conditions: (1) the image finiteness of transition steps and (2) the existence of *a pair of distinguishing formulae* for logically inequivalent procceses. The former is usually assumed by definition of the labeled transition system and the latter would usually be obliged to be derived using the former.

The first condition, image finiteness, means that a process can only have finitely many next step process configurations for each label. Assuming image finiteness for labeled transition systems are fairly standard. For instance, it is also assumed in proofs for the adequacy of Hennessy–Milner logic (HML) for CCS [10]. In our context of proving the closedness lemma `sateq_one_L`, image finiteness amounts to the existence of finitely many Q_is such that $[\![Q \xrightarrow{\alpha} Q_i]\!]$. Assuming this, we can exhaustively enumerate all such instances as $Q_1, Q_2, \cdots, Q_m$.

The remaining proof for `sateq_one_L` is to show that at least one of $P_1 \overset{\models}{\sim} Q_1, \cdots, P_1 \overset{\models}{\sim} Q_m$ holds, that is, it is impossible to fail all of them. Let us assume that it all fail, that is, $P_1 \overset{\models}{\not\sim} Q_1, \cdots, P_1 \overset{\models}{\not\sim} Q_m$. If we can find a formula ϕ_i for each $P_1 \overset{\models}{\not\sim} Q_i$ such that $P_1 \models \phi_i$ holds but $Q_i \models \phi_i$ does not hold, then it contradicts the assumption $P \overset{\models}{\sim} Q$ because we can build a formula $\phi = \langle \alpha \rangle \bigwedge_{i=1\ldots m} \phi_i$, which makes $P \models \phi$ hold but not $Q \models \phi$. Thus, the essence of the closedness lemma proof is whether we can always find a formula ϕ_i such that $P_1 \models \phi_i$ holds but $Q_i \models \phi_i$ does not hold, given that $P_1 \overset{\models}{\not\sim} Q_i$.

Having the second condition below guarantees that we can always find such ϕ_i. There are two reasons for $P_1 \overset{\models}{\not\sim} Q_i$:

- $\exists \phi_i$ such that $P_1 \models \phi_i$ holds but $Q_i \models \phi_i$ does not.
- $\exists \phi_i'$ such that $Q_i \models \phi_i'$ holds but $P_1 \models \phi_i'$ does not.

Thus, $P_1 \overset{\models}{\not\sim} Q_i$ guarantees the existence of either ϕ_i or ϕ_i'. However, the existence of one of them may not necessarily guarantee the other. In cases of HML for CCS and MPW for late bisimulation in the π-calculus we know that both ϕ_i and ϕ_i' exists because the negation of formulae can be defined syntactically. However, the encodings of the input box and diamond modalities for $\mathcal{OM}$ ($\Box^{\downarrow}$and $\Diamond^{\downarrow}$ in Fig. 2) are not logically dual; only if one of the $\forall$ `w` in lines 28 and 29 were $\exists$ `w` could the modalities be classically dual to each

other. Therefore, syntactic negation for $\mathcal{OM}$ formulas cannot be defined in general, even before taking free variables and the intuitionistic framework into consideration.

We would need even more careful argument to establish the existence of the pair of formulae ϕ_i and ϕ_i' in the presence of free variables. That is, we need to establish the following: for any processes P and Q such that $P \not\sim Q$ fails, that is $P \not\nsim Q$, we can find a pair of formulae that is satisfied by P or Q respectively but fails for the other process, despite the absence of syntactic negation. Our key observation in establishing this is to use open bisimulation instead of logical equivalence. From ($\sim \subseteq \not\sim$), which is established in Section III-A, we get ($\not\nsim \subseteq \nsim$) by de Morgan's law. Therefore, it suffices to find a pair of formulae for $P \nsim Q$, instead of $P \not\nsim Q$. In Section IV, we show how to find this pair of distinguishing formulae (Proposition 4.8).

C. Encoding of the history in direct semantics

So far we have only considered universal quantification, which encodes free variables and also inputs. Free variables are modeled as inputs from the top-level environment both in direct semantics and in Abella encoding. It is sufficient to consider universal quantifications only while formalizing the soundness and completeness theorems. More generally, during the internal steps of bisimulation and satisfaction evaluation, fresh names from bounded outputs may be generated. These names are encoded as ∇-quantified variables in Abella.

For instance, $[\![F(x,y,z) \models^{x^i y^o z^i} P(x,y,z)]\!] \triangleq$

$$\forall x, \nabla y, \forall z,\ \texttt{sat}\ ([\![F]\!]\, x\, y\, z)\ ([\![P]\!]\, x\, y\, z).$$

∇-quantified variables are distinct from the other variables already in scope, hence, $x = y$ is false. However, $y = z$ is not necessarily false because z is introduced afterwards so that there exists a substitution $\theta = \{^{y}/_{z}\}$ such that $y\theta = z\theta$. This exactly coincides with the restriction on respectful substituions due to bound outputs in histories according to Def. 2.1.

In Section V-B, we discuss related work on this style of encoding the π-calculus and its properties within a logic that supports higher-order abstract syntax and ∇-quantifiers.

IV. Completeness of Open Bisimulation

There is a constructive definition of non-bisimulation. Since bisimulation is defined in terms of a greatest fixed point of relations satisfying a certain closure property, non-bisimulation is defined in terms of a least fixed point satisfying the dual property. This leads to the following constructive definition of non-bisimilarity from which a non-bisimulation algorithm can be extracted. Since non-bisimilarity is defined in terms of a least fixed point, there is a finite winning strategy, consisting of a finite tree of moves such that in each branch eventually a pair of processes and a history is reached such that one process can make a move that the other cannot always match.

Definition 4.1 (non-bisimulation): Firstly, we inductively define the family of relation $\nsim_n^h$, for $n \in \mathbb{N}$. The base case is when, for some respectful substitution one player can make a move, that cannot be matched by the other player without assuming a stronger substitution. The class of all such pairs of processes form the base case for the construction of the non-bisimilarity relation, say $P \nsim_0^h Q$. More precisely, the relation $\nsim_0^h$ is the least symmetric relation such that for any P and Q, $P \nsim_0^h Q$ whenever there exist process P', action π and substitution σ respecting h, such that the following holds.

- $P\sigma \xrightarrow{\pi} P'$ and there is no Q' such that $Q\sigma \xrightarrow{\pi} Q'$.

Inductively, $\nsim_{n+1}^h$ is the least symmetric relation extending $\nsim_n^h$ such that $P \nsim_{n+1}^h Q$ whenever for some substitution σ respecting h, at least one of the following holds, where α is τ or $\overline{a}b$:

- $P\sigma \xrightarrow{\alpha} P'$ and for all Q_i such that $Q\sigma \xrightarrow{\alpha} Q_i$, $P' \nsim_n^{h\sigma} Q_i$.
- $P\sigma \xrightarrow{\overline{a}(x)} P'$, where x is fresh for $P\sigma$, $Q\sigma$ and $h\sigma$, and for all Q_i such that $Q\sigma \xrightarrow{\overline{a}(x)} Q_i$, $P' \nsim_n^{h\sigma \cdot x^o} Q_i$.
- $P\sigma \xrightarrow{a(x)} P'$, where x is fresh for $P\sigma$, $Q\sigma$ and $h\sigma$, and for all Q_i such that $Q\sigma \xrightarrow{a(x)} Q_i$, $P' \nsim_n^{h\sigma \cdot x^i} Q_i$.

Thereby, the relation $P \nsim_n^h Q$ contains all processes that can be distinguished by a strategy with depth at most n, i.e. at most n moves are required to reach a pair of processes in $\nsim_0^h$, at which point there is an accessible world in which a process can make a move that the other process cannot match.

The relation $\nsim^h$, pronounced non-bisimilarity with history h, is defined to be the least relation containing $\nsim_n^h$ for all $n \in \mathbb{N}$, i.e. $\bigcup_{n\in\mathbb{N}} \nsim_n^h$.

h) Preliminaries: We required the following terminology for substitutions, and abbreviations for formulae.

Definition 4.2: Composition of substitutions σ and θ is defined such that $P(\sigma \cdot \theta) = (P\sigma)\theta$, for all processes P. For substitutions σ and θ, $\sigma \leq \theta$ whenever there exists σ' such that $\sigma \cdot \sigma' = \theta$.

For a finite substitution $\sigma = \{^{z_1}/_{x_1}\} \cdots \{^{z_n}/_{x_n}\}$ the formula $[\sigma]\phi$ abbreviates $[x_1 = z_1] \ldots [x_n = z_n]\phi$. For a finite set of formulae ϕ_i, the formulae $\bigvee_i \phi_i$ abbreviates $\phi_1 \vee \ldots \vee \phi_n$, where the empty disjunction is $\mathbb{ff}$. Similarly $\bigwedge_i \phi_i$ abbreviates $\phi_1 \wedge \ldots \wedge \phi_n$, where the empty conjunction is $\mathbb{tt}$.

We require the following technical lemmas. The first ensures there are finitely many equalities between names characterising all substitutions that enable a transition. The second (image finiteness) ensures that there are finitely many reachable states in one step, up to renaming. The third extends the definition of the box-match modality to finite substitutions. The fourth is required in inductive cases involving bound output and input. The fifth is a monotonicity property for satisfaction.

Lemma 4.3: For all P there exists finitely many σ_j such that for all σ, if $P\sigma \xrightarrow{\alpha} Q$, then there exists j and σ', such that $\sigma_j \cdot \sigma' = \sigma$ and $P\sigma_j \xrightarrow{\alpha} Q'$ such that $Q'\sigma' = Q$.

Lemma 4.4: For process P and action π there are finitely many P_i such that $P \xrightarrow{\pi} P_i$, up to α-conversion.

Lemma 4.5: If for all θ such that θ respects h and $\sigma \leq \theta$, it holds that $P\theta \models^{h\theta} \phi\theta$, then $P \models^h [\sigma]\phi$ holds.

Lemma 4.6: If $\sigma \cdot \theta$ respects h, then θ respects $h\sigma$.

Lemma 4.7: If θ respects h and $P \models^h \phi$, then $P\theta \models^{h\theta} \phi\theta$.

i) Algorithm for distinguishing formulae: The constructive definition of non-bisimulation gives a tree of substitutions and actions forming a strategy showing that two processes are not bisimilar. The following proposition shows that $\mathcal{OM}$ formulae are sufficient to capture such strategies. For any strategy that distinguishes two processes, we can construct *distinguishing* $\mathcal{OM}$ formulae. A distinguishing formula holds for one process but not for the other process. Furthermore, there are always at least two distinguishing formulae, one biased to the left and the other biased to the right, as in the construction of the proof for the following proposition. As we do not have the excluded middle in our logic, the left bias cannot be simply obtained by negating the right bias and vice versa; both need to be constructed simultaneously and they may not be related by negation.

Proposition 4.8: If $P \nsim Q$ then there exists ϕ_L such that $P \models \phi_L$ and $Q \not\models \phi_L$, and also there exists ϕ_R such that $Q \models \phi_R$ and $P \not\models \phi_R$.

Proof: Since $\nsim^h$ is defined by a least fixed point over a family or relations $\nsim^h_n$, if $P \nsim^h Q$, there exists n such that $P \nsim^h_n Q$, so we can proceed by induction on the depth of a winning strategy.

In the base case, assume that $P \nsim^h_0 Q$, hence by definition, for substitution σ respecting h, $P\sigma \xrightarrow{\pi} P'$, such that there is no Q' such that $Q\sigma \xrightarrow{\pi} Q'$, up to symmetry of $\nsim^h_n$. By Lemma 4.3, there exist finitely many pairs of names x_j and y_j such that $x_j\sigma \neq y_j\sigma$, and, for any substitution θ, if $Q\theta \xrightarrow{\pi} Q'$ there exists j such that $x_j\theta = y_j\theta$. In this case, there are two distinguishing formulae $[\sigma]\langle\pi\rangle\mathtt{tt}$ and $[\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$ biased to P and Q respectively. There are four cases to check to determine that these are indeed distinguishing formulae.

Case $P \models^h [\sigma]\langle\pi\rangle\mathtt{tt}$**:** Consider all θ such that $\sigma \leq \theta$ and θ respects h. By definition there exists θ' such that $\sigma \cdot \theta' = \theta$, so since $P\sigma \xrightarrow{\pi} P'$, $P\theta \xrightarrow{\pi} P'\theta'$. Thereby, since $P'\theta' \models^{h'} \mathtt{tt}$ holds, $P\theta \models^h \langle\pi\rangle\mathtt{tt}$. Since θ respects h and $\sigma \leq \theta$, by Lemma 4.5, $P \models^h [\sigma]\langle\pi\rangle\mathtt{tt}$ holds.

Case $Q \not\models^h [\sigma]\langle\pi\rangle\mathtt{tt}$**:** Assume $Q \models^h [\sigma]\langle\pi\rangle\mathtt{tt}$ for contradiction. Now, since σ respects h and $\sigma \leq \sigma$, by Lemma 4.5, if $Q \models^h [\sigma]\langle\pi\rangle\mathtt{tt}$ holds then $Q\sigma \models^h \langle\pi\rangle\mathtt{tt}$ must hold; which holds only if there exists Q' such that $Q\sigma \xrightarrow{\pi} Q'$, contradicting the assumption that no such Q' exists. Thereby $Q \not\models^h [\sigma]\langle\pi\rangle\mathtt{tt}$.

Case $Q \models^h [\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$**:** Consider all substitutions θ respecting h and Q' such that $Q\theta \xrightarrow{\pi} Q'$. It must be the case that there exists j such that $x_j\theta = y_j\theta$, thereby it holds that $Q' \models^h \langle x_j\theta = y_j\theta\rangle\mathtt{tt}\theta$; hence clearly $Q' \models \bigvee_j\langle x_j = y_j\rangle\mathtt{tt}\theta$ holds. Hence we conclude $Q \models^h [\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$.

Case $P \not\models^h [\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$**:** Assume for contradiction $P \models^h [\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$. This holds iff for all processes S and substitutions θ respecting h, $P\theta \xrightarrow{\pi} S$ implies $S \models^{h'} \bigvee_j\langle x_j = y_j\rangle\mathtt{tt}\theta$. Since we know that σ respects h and $P\sigma \xrightarrow{\pi} P'$, we have $P' \models^{h'} \bigvee_j\langle x_j = y_j\rangle\mathtt{tt}\sigma$. This holds only if for some j, $P' \models^{h'} \langle x_j = y_j\rangle\mathtt{tt}\sigma$; hence, $x_j\sigma = y_j\sigma$ for some j, which contradicts the assumption that $x_j\sigma \neq y_j\sigma$. Thereby $P \not\models^h [\pi]\bigvee_j\langle x_j = y_j\rangle\mathtt{tt}$.

Now consider the inductive case. Given P, Q, if $P \nsim^h_{n+1} Q$, up to symmetry of $\nsim^h_{n+1}$, there are three cases to consider, for some substitution σ respecting h, where α is either τ or $\overline{a}b$:

- $P\sigma \xrightarrow{\alpha} P'$ and for all Q_i such that $Q\sigma \xrightarrow{\alpha} Q_i$, $P' \nsim^{h\sigma}_n Q_i$.
- $P\sigma \xrightarrow{\overline{a}(x)} P'$, where x is fresh for $P\sigma$, $Q\sigma$ and $h\sigma$, and for all Q_i such that $Q\sigma \xrightarrow{\overline{a}(x)} Q_i$, $P' \nsim^{h\sigma\cdot x^o}_n Q_i$.
- $P\sigma \xrightarrow{a(x)} P'$, where x is fresh for $P\sigma$, $Q\sigma$ and $h\sigma$, and for all Q_i such that $Q\sigma \xrightarrow{a(x)} Q_i$, $P' \nsim^{h\sigma\cdot x^i}_n Q_i$.

We consider the second case above involving bound output only, the other two cases are similar — differing only in the accounting for respectful substitutions according to Def. 2.1.

For $P\sigma \xrightarrow{\overline{a}(x)} P'$, by Lemma 4.4, there exist finitely many Q_i such that $Q\sigma \xrightarrow{\overline{a}(x)} Q_i$. Since, for each i, $P' \nsim^h_n Q_i$, by the induction hypothesis, for all i, since $P' \nsim^{h\sigma\cdot x^o}_n Q_i$ there exist ϕ^L_i and ϕ^R_i such that $P' \models^{h\sigma\cdot x^o} \phi^L_i$ and $Q_i \not\models^{h\sigma\cdot x^o} \phi^L_i$ and $P' \not\models^{h\sigma\cdot x^o} \phi^R_i$ and $Q_i \models^{h\sigma\cdot x^o} \phi^R_i$. Furthermore, assume that σ is minimal with respect to the order over substitutions in the sense that, if $\theta \leq \sigma$ is such that $P\theta \xrightarrow{\overline{a}(x)} P''$, where x is fresh for $P\theta$, $Q\theta$ and $h\theta$, and for all Q'' such that $Q\theta \xrightarrow{\overline{a}(x)} Q''$, $P'' \nsim^{h\theta\cdot x^o}_n Q''$, then $\theta = \sigma$. By Lemma 4.3, there are finitely many variables x_j and y_j such that $x_j\sigma \neq y_j\sigma$ and for any substitution θ such that, for some S, $Q\theta \xrightarrow{\overline{a}(x)} S$ either: there exist j such that $x_j\theta = y_j\theta$; or $\sigma \leq \theta$ and there exists θ' such that $Q\sigma \xrightarrow{\overline{a}(x)} S'$ and $S'\theta' = S$. Notice that cases where $\theta \leq \sigma$ are eliminated by the minimality of σ. From the above we can construct the distinguishing formulae $[\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$ and $[\overline{a}(x)](\bigvee_i\phi^R_i \vee \bigvee_j\langle x_j = y_j\rangle)$. There are four cases to consider to verify these are distinguishing formulae.

Case $P \models^h [\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$**:** Consider all θ such that $\sigma \leq \theta$, θ respects h, and without loss of generality x is fresh such that $x \notin \mathrm{dom}(\theta)$ and $x \notin \mathrm{fv}(h\theta)$. By definition, there exists θ' such that $\sigma \cdot \theta' = \theta$. Now since $\sigma \cdot \theta'$ respects h, by Lemma 4.6, θ' respects $h\sigma$ hence since $x \notin \mathrm{dom}(\theta')$ and $x \notin \mathrm{fv}(h\sigma\theta')$, θ' respects $h\sigma\cdot x^o$. Thereby since θ' respects $h\sigma\cdot x^o$ and $P' \models^{h\sigma\cdot x^o} \phi^L_i$ holds, by Lemma 4.7, $P'\theta' \models^{h\theta\cdot x^o} \phi^L_i\theta'$ holds. The above holds for all i, hence $P'\theta' \models^{h\theta\cdot x^o} \bigwedge_i\phi^L_i\theta'$ holds. Now, since $P\sigma \xrightarrow{\overline{a}(x)} P'$, $P\theta \xrightarrow{\overline{a}(x)} P'\theta'$ and hence $P\theta \models^{h\theta} \langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i\theta'$ holds. Thereby, by Lemma 4.5, since $\sigma \leq \theta$ and θ respects h, $P \models^h [\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$ holds, as required.

Case $Q \not\models^h [\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$**:** Assume for contradiction that $Q \models^h [\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$. Since σ respects h and trivially $\sigma \leq \sigma$, by Lemma 4.5, the above assumption holds only if $Q\sigma \models^{h\sigma} \langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$ holds. Now $Q\sigma \models^{h\sigma} \langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$ holds only if there exists Q' such that $Q\sigma \xrightarrow{\overline{a}(x)} Q'$ and $Q' \models^{h\sigma\cdot x^o} \bigwedge_i\phi^L_i$, which holds only if for all i, $Q' \models^{h\sigma\cdot x^o} \phi^L_i$. Notice that $Q' = Q_k$ for some k, and therefore $Q_k \models^{h\sigma\cdot x^o} \phi^L_k$; but it was assumed that $Q_k \not\models^{h\sigma\cdot x^o} \phi^L_k$ leading to a contradiction. Therefore $Q \not\models^h [\sigma]\langle\overline{a}(x)\rangle\bigwedge_i\phi^L_i$.

Case $Q \models [\overline{a}(x)](\bigvee_i\phi^R_i \vee \bigvee_j\langle x_j = y_j\rangle\mathtt{tt})$**:** Consider all Q' and θ respecting h such that $Q\theta \xrightarrow{\overline{a}(x)} Q'$ and without loss of generality x is fresh such that $x \notin \mathrm{dom}(\theta)$ and $x \notin \mathrm{fv}(h\theta)$. There are two sub-cases to consider. Firstly consider where for some k, $x_k\theta = y_k\theta$, in which case $Q' \models^{h\theta\cdot x^o} \langle x_k\theta = y_k\theta\rangle\mathtt{tt}$, hence $Q' \models^{h\theta\cdot x^o} (\bigvee_i\phi^R_i \vee \bigvee_j\langle x_j = y_j\rangle\mathtt{tt})\theta$, by definition of

disjunction. Secondly consider where there exists θ' such that $\sigma \cdot \theta' = \theta$ and for some ℓ, $Q\sigma \xrightarrow{\overline{a}(x)} Q_\ell$ such that $Q_\ell\theta' = Q'$. Now since $\sigma \cdot \theta'$ respects h, by Lemma 4.6, θ' respects $h\sigma$, hence by definition of respectful substitutions, since $x \notin \mathrm{dom}(\theta)$ and $x \notin \mathrm{fv}(h\theta)$, θ' respects $h\sigma \cdot x^o$. Thereby, by Lemma 4.7, since $Q_\ell \models^{h\sigma\cdot x^o} \phi^R_\ell$ and θ' respects $h\sigma \cdot x^o$, $Q_\ell\theta' \models^{h\theta\cdot x^o} \phi^R_\ell\theta'$. Note that $\phi^R_\ell\sigma = \phi^R_\ell$ hence $\phi^R_\ell\theta = \phi^R_\ell\theta'$, and therefore $Q' \models^{h\theta\cdot x^o} \phi^R_\ell\theta$. Thus, by definition of disjunction, $Q' \models^{h\theta\cdot x^o} (\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})\theta$ holds, as required. Thus by definition of $[\overline{a}(x)]$, we can conclude that $Q \models^h [\overline{a}(x)](\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})$ holds.

Case $P \not\models^h [\overline{a}(x)](\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})$**:** Assume that for contradiction $P \models^h [\overline{a}(x)](\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})$. Since σ respects h and $P\sigma \xrightarrow{\overline{a}(x)} P'$, the assumption holds only if $P' \models^{h\sigma\cdot x^o} (\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})\sigma$. This holds only if, for some i, $P' \models^{h\sigma\cdot x^o} \phi^R_i$, or, for some j, $P' \models^{h\sigma\cdot x^o} \langle x_j\sigma = y_j\sigma\rangle \mathtt{tt}$. However, for all i, $P' \not\models \phi^R_i$; and also, since for all j, $x_j\sigma \neq y_j\sigma$, $P' \not\models \langle x_j\sigma = y_j\sigma\rangle \mathtt{tt}$, leading to a contradiction in either case. Thereby $P \not\models^h [\overline{a}(x)](\bigvee_i \phi^R_i \vee \bigvee_j \langle x_j = y_j\rangle \mathtt{tt})$.

By induction we have established that for any history h, processes P and Q, and any n, if $P \not\sim^h_n Q$ then there exists ϕ_L such that $P \models^h \phi_L$ and $Q \not\models^h \phi_L$, and also there exists ϕ_R such that $Q \models^h \phi_R$ and $P \not\models^h \phi_R$. The result then follows by observing that $\not\sim^h$ is least relation containing all $\not\sim^h_n$; and, furthermore, $P \not\sim Q$ holds simply when $P \not\sim^{x^i_1\cdots x^i_n} Q$ holds, where $\mathrm{fv}(P) \cup \mathrm{fv}(Q) \subseteq \{x^i_1, \ldots, x^i_n\}$. ■

Since open bisimulation is decidable for finite π-calculus processes, the constructive non-bisimulation in Definition 4.1 coincides with the negation of open bisimulation.

Lemma 4.9: For finite processes, $P \not\sim^h Q$ holds, according to constructive non-bisimulation in Definition 4.1, if and only if $P \sim Q$ does not hold.

Combining Proposition 4.8 with Lemma 4.9 yields immediately the completeness of $\mathcal{OM}$ with respect to open bisimulation. Completeness (Theorem 2.6) establishes that the set of all pairs of processes that have the same set of distinguishing formulae is a bisimulation. The proof can now be stated as follows.

Proof of Theorem 2.6: Assume that for finite processes P and Q, for all formulae ϕ, $P \models \phi$ iff $Q \models \phi$. Now for contradiction suppose that $P \sim Q$ does not hold. By Lemma 4.9, $P \not\sim Q$ must hold. Hence by Proposition 4.8 there exists ϕ_L such that $P \models \phi_L$ but $Q \not\models \phi_L$, but by the assumption above $Q \models \phi_L$, leading to a contradiction. Thereby $P \sim Q$. ■

Notice that soundness (Theorem 2.5) and the non-bisimulation algorithm (Proposition 4.8) also hold for infinite π-calculus processes (defined using replication for instance). However, for infinite π-calculus processes with replication open bisimulation is undecidable; hence Lemma 4.9 cannot be assumed. Thereby in the infinite case, a weaker result than completeness holds. While it is impossible that $P \sim Q$ and $P \not\sim Q$ holds, it may be the case that neither holds; hence a more expressive logic is required to completely characterise open bisimulation in the case of infinite processes. Future work includes investigating extensions of $\mathcal{OM}$ with fixed points.

A. Example Runs of Distinguishing Formula Algorithm

We provide further examples of non-bisimilar processes that illustrate subtle aspects of the algorithm. In particular, these examples illustrate the need for disjunctions of postconditions in both the base case and inductive steps.

j) Processes with multiple postconditions: Consider two non-bisimilar processes $[x = y]\tau + [w = z]\tau \not\sim \tau$. Observe that $\tau \xrightarrow{\tau} 0$ but $([x = y]\tau + [w = z]\tau)\theta \xrightarrow{\tau}$ only if it is either $x\theta = y\theta$ or $w\theta = z\theta$, or both. Thus, we have the distinguishing formula for the left process as below

$$[x = y]\tau + [w = z]\tau \models [\tau](\langle x = y\rangle \mathtt{tt} \vee \langle w = z\rangle \mathtt{tt})$$

and for the right simply $\tau \models \langle\tau\rangle \mathtt{tt}$.

k) Processes with postconditions in inductive case: Firstly observe that $\overline{a}a + \overline{b}b \not\sim \overline{a}a$ are distinguished since $\overline{a}a + \overline{b}b \xrightarrow{\overline{b}b} 0$, by $\overline{a}a$ cannot make a $\overline{b}b$ transition under any substitution. Hence we have the distinguishing formulae $\overline{a}a + \overline{b}b \models \langle\overline{b}b\rangle \mathtt{tt}$ and $\overline{a}a \models [\overline{b}b]\mathtt{ff}$,

Inductively, consider the following non-bisimilar processes.

$$P \triangleq \tau.(\overline{a}a + \overline{b}b) + [x = y]\tau.\overline{a}a \not\sim \tau.(\overline{a}a + \overline{b}b) + \tau.\overline{a}a \triangleq Q$$

This inductive case is lead by $Q \xrightarrow{\tau} \overline{a}a$ and which can only be matched by $P \xrightarrow{\tau} \overline{a}a + \overline{b}b$. By the above observation, we have distinguishing formulae for $\overline{a}a + \overline{b}b \not\sim \overline{a}a$. Furthermore, $P\theta \xrightarrow{\tau}$ for substitutions θ such that $x\theta = y\theta$.

This leads to the following distinguishing formula for the left side, consisting of a box τ followed by a disjunction of the left distinguishing formula for $\overline{a}a + \overline{b}b \not\sim \overline{a}a$, and the postcondition for any additional τ transitions.

$$\tau.(\overline{a}a + \overline{b}b) + [x = y]\tau.\overline{a}a \models [\tau](\langle\overline{b}b\rangle \mathtt{tt} \vee \langle x = y\rangle \mathtt{tt})$$

The distinguishing formula for the right process is diamond τ followed by the right distinguishing formula for $\overline{a}a + \overline{b}b \not\sim \overline{a}a$, as follows: $\tau.(\overline{a}a + \overline{b}b) + \tau.\overline{a}a \models \langle\tau\rangle[\overline{b}b]\mathtt{ff}$.

l) Further examples: We briefly mention two more examples. The following processes are not bisimilar:

$$[x = y]\tau.\tau + \tau \quad \not\sim \quad \tau.\tau + \tau$$

For a distinguishing formula biased to the left process, by "thinking" we obtain the following distinguishing formula.

$$[x = y]\tau.\tau + \tau \models [\tau][\tau]\langle x = y\rangle$$

However, the algorithm will generate an alternative distinguishing formula, also correctly biased to the left process.

$$[x = y]\tau.\tau + \tau \models [\tau]([\tau]\mathtt{ff} \vee \langle x = y\rangle \mathtt{tt})$$

Thus a human may be able to find a more concise distinguishing formula than the algorithm.

For a more elaborate example consider the following.

$$P \triangleq \tau.(\overbrace{\tau + \tau.[x = y]\tau}^{P'}) + \tau.(\tau + \tau.[x = y][w = z]\tau) \not\sim \tau.(\underbrace{\tau + \tau.[x = y][w = z]\tau}_{Q'}) \triangleq Q$$

The non-bisimulation strategy involves two switches of the leading side: firstly, lead by $P \xrightarrow{\tau} P'$ on the left, matched by $Q \xrightarrow{\tau} Q'$; secondly, lead by $Q' \xrightarrow{\tau} [x = y][w = z]\tau$ on the right, matched by $P' \xrightarrow{\tau} 0$ and $P' \xrightarrow{\tau} [x = y]\tau$; thirdly, lead by

$[x = y]\tau\{^{y}/_{x}\} \xrightarrow{\tau} 0$ again on the left, which cannot be matched by the other side $[x = y][w = z]\tau\{^{y}/_{x}\}$.

The algorithm in the completeness proof generates the following distinguishing formulae:

$$P \models \langle\tau\rangle[\tau]([\tau]\mathit{ff} \vee [x = y]\langle\tau\rangle\mathit{tt}) \quad Q \models [\tau]\langle\tau\rangle[\tau]\langle w = z\rangle$$

Note if we assume the law of excluded middle, both processes above would be equivalent to $\tau.(\tau+\tau.\tau)$. Fortunately, we do not assume the law of excluded middle.

V. Related Work

We consider the relationship between the intuitionistic modal logic for open bisimulation presented in this work and established classical logics. We also compare this work to existing work claiming to characterise open bisimulation for the π-calculus.

A. Comparison to Classical Logics for Late Bisimulation

The late Milner-Parrow-Walker logic, called $\mathcal{LM}$ in the literature [4] for "($\mathcal{L}$) late modality with ($\mathcal{M}$) match" differs from the logic presented in this paper in three significant ways: Firstly, free names are a priori assumed to be distinct; secondly, we assume the law of excluded middle of classical logic, and thirdly the late input box modality is defined differently as follows:

- $P \models^{L} [a(x)]^{L}\phi$ iff for all Q such that $P \xrightarrow{a(x)} Q$ there exists σ such that $Q \models^{L} \phi\sigma$.

To see that $\mathcal{LM}$ is not a congruence consider processes $[x = y]\overline{x}x$ and 0. These processes satisfy the same set of late formulae (any formula equivalent to tt), since, for $\mathcal{LM}$, x and y are a priori assumed to be distinct names. However, $a(y).[x = y]\overline{x}x$ and $a(y).0$ have distinguishing formulae $a(y).[x = y]\overline{x}x \models [a(y)]^{L}\langle\overline{x}x\rangle\mathit{tt}$ and its de Morgan complement $a(y).0 \models \langle a(y)\rangle^{L}[\overline{x}x]\mathit{ff}$. In the former case there exists substitution $\{^{x}/_{y}\}$ such that $[x = y]\overline{x}x\{^{x}/_{y}\} \models \langle\overline{x}x\rangle\mathit{tt}\{^{x}/_{y}\}$. In the later case, 0 cannot ever perform a $\overline{x}x$-transition, hence $0 \models [\overline{x}x]\mathit{ff}$ holds.

Between open bisimulation and late bisimulation there is late congruence, which is the greatest congruence relation contained in late bisimulation. Late congruence must contain open bisimilarity, since open bisimulation is contained in late bisimulation and open bisimilarity is a congruence. Late congruence also has a simpler characterisation: P and Q are late congruent whenever for all substitutions σ, and $P\sigma$ is late bisimilar to $Q\sigma$. The quantification over all substitutions, combined with the law of excluded middle, has the effect that we check late bisimulation with respect to all combinations of equalities and inequalities between free names.

As for open bisimulation, $[x = y]\overline{x}x$ and 0 are not late congruent. This is because for substitution $\{^{x}/_{y}\}$, $([x = y]\overline{x}x)\{^{x}/_{y}\}$ and $0\{^{x}/_{y}\}$ are clearly not late bisimilar. This illustrates that late congruence is strictly finer than late bisimulation. However, open bisimulation is still strictly finer than late congruence, since $\tau + \tau.\tau + \tau.[x = y]\tau$ and $\tau + \tau.\tau$ are late congruent. Late congruence holds since, for any substitution θ, $(\tau + \tau.\tau)\theta$ and $(\tau + \tau.\tau + \tau.[x = y]\tau)\theta$ are late bisimilar. In contrast, we know these processes are not open bisimilar; and furthermore, have distinguishing formula that rely on the absence of the law of excluded middle.

B. Embeddings into an intuitionistic logic with λ-syntax (HOAS), induction, nabla (∇), and coinduction

Tiu and Miller [5] studied shallow embeddings of the π-calculus into the logic LINC, as well as late and open bisimulation and their respectful modal logics. This is the most closely related work to ours because our encodings in Abella (Section III) were adapted from their work.

In their encoding, both late and open bisimulations are encoded by essentially the same modalities, differing only in the the law of the excluded middle over names and universal quantification of free variables. However, neither examples of distinguishing formulae for open bisimulation nor the observation regarding the absence of syntactic negation, which is crucial for establishing completeness for the logic characterizing open bisimulation, was provided.

The main departure from the approach in this work, is Tiu and Miller's reuse of the input box modality from $\mathcal{LM}$, which involves an existential quantifaction over substitutions. In contrast, our input box modality in $\mathcal{OM}$ involves universal quantification over all respectful substitutions. Our choice in $\mathcal{OM}$ is critical for generating distinguishing formulae. For example, the following processes are not open bisimilar.

$$a(x).\tau + a(x) + a(x).[x = a]\tau \quad \not\sim \quad a(x).\tau + a(x)$$

For the above processes, the algorithm for distinguishing formulae in Proposition 4.8, correctly generates the following $\mathcal{OM}$ formula biased to the right.

$$a(x).\tau + a(x) \models [a(x)]\,(\langle\tau\rangle\mathit{tt} \vee [\tau]\mathit{ff})$$

However, using only late modalities, there is no distinguishing formula for these processes biased to the right: e.g., the formula with a late modality $[a(x)]^{L}\,(\langle\tau\rangle\mathit{tt} \vee [\tau]\mathit{ff})$ succeeds for both processes, even when rejecting the law of excluded middle. Furthermore, the choice of modalities we make in $\mathcal{OM}$ is intuitively correct, since in open bisimulation the choice of selecting a substitution is deferred as late as possible — possibly several transitions later.

C. A generic formalization using nominal logic

Recently, Parrow et al. [11] provided a general proof of the soundness and completeness of logical equivalence for various modal logics with respect to corresponding bisimulations. The proof is parametric on properties of substitutions, which can be instantiated for a range of bisimulations. Moreover, their proof is mechanised using Nominal Isabelle. The conference version [11], sketches how to instantiate the abstract framework for open bisimulation in the π-calculus without input prefixes only. However, we understand from communication with the authors that open bisimulation for the π-calculus with input prefixes will be covered in a forthcoming extended version.

Stylistically, our intuitionistic modal logic is quite different from an instantiation of the abstract framework of Parrow et al. for open bisimulation. Their framework, is classical

and works by syntactically restricting "effect" modalities in formulae, depending on the type of bisimulation. Their effects represent substitutions that reach worlds permitted by the type of bisimulation. In contrast, the modalities of the intuitionistic modal logic $\mathcal{OM}$ in this paper are syntactically closer to long established modalities for the π-calculus [4]; differing instead in their semantic interpretation and in the absence of classical negation. An explanation for the stylistic differences is that for every intuitionistic logic, such as the intuitionistic modal logic in this work, there should be a corresponding classical modal logic based on an underlying Kripke semantics. Such a Kripke semantics would reflect the accessible worlds, as achieved by the syntactically restricted effect modalities in the abstract classical framework instantiated for open bisimulation.

D. Modal Logics and Open Bisimulation for the Spi-Calculus

Frendrup and others [12] presented modal logics that characterise early and late bisimulations for the spi-calculus [13], which is a cryptographic extension to the π-calculus. Afterwards, Tui [14] defined an open bisimulation for the spi-calculus based on bi-traces. Finding a modal logic that characterise open bisimulation for spi-calculus is an open research problem, however we believe techniques and definitions in this work extend to the spi-calculus. Histories in our semantics are in fact inspired by the bi-traces used to define open bisimulation for the spi-calculus [14]; all histories in this work correspond to bi-traces for the spi-calculus where messages only involve names and no other cryptographic primitives.

VI. Conclusion and Future Work

The main result of this paper is a sound and complete logical characterisation of open bisimulation for the π-calculus. To achieve this result, we introduce modal logic $\mathcal{OM}$, defined in Fig. 1. The soundness of $\mathcal{OM}$ with respect to open bisimulation, Theorem 2.5, is mechanically proven in Abella as explained in Section III. The details of the completeness, Theorem 2.6, are provided in Section IV.

There are several novel features of $\mathcal{OM}$ compared to established modal logics for π-calculus, such as $\mathcal{LM}$ characterising late bisimulation. Firstly, as demonstrated in Examples II-B0e and IV-A0l, the absence of the law of excluded middle is essential for the existence of distinguishing formulae in $\mathcal{OM}$ for certain processes that are not open bisimilar (but are late congruent). The absence of the law of excluded middle is an intuitionistic assumption; and, as explained in the introduction, $\mathcal{OM}$ can indeed be considered to be a conservative extension of intuitionistic logic. Furthermore, in contrast to classical modal logics such as $\mathcal{LM}$, diamond and box modalities have independent interpretations, not dual to each other. These properties are expected under criterion set out for intuitionistic modal logics [15]. The absence of de Morgan dualities complicates the construction of distinguishing formulae.

The completeness proof involves an algorithm, Proposition 4.8, that constructs distinguishing formulae for non-bisimilar processes. To use this algorithm, firstly attempt to prove that two processes are open bisimilar. If they are non-bisimilar, after a finite number of steps, a distinguishing strategy, according to Def. 4.1, will be discovered. The strategy can then be used to inductively construct two distinguishing formulae, biased to each process. A key feature of the construction is that there are restricted versions of absolute truth by preconditions ($[\sigma]\langle\pi\rangle\mathtt{tt}$ restricted from $\langle\pi\rangle\mathtt{tt}$) and, dually, there are relaxed versions of absolute falsity by post-conditions ($[\pi]\langle\sigma\rangle\mathtt{tt}$ relaxed from $[\pi]\mathtt{ff}$), as demonstrated in Examples II-A0d, IV-A0j and IV-A0k.

Our logic $\mathcal{OM}$ is suitable for formal and automated reasoning; in particular, it has natural encodings in Abella for mechanised reasoning, demonstrated in Section III, and Bedwyr [16] for automatic proof search. All the bisimulation and satisfaction judgements in examples have been automatically checked in Bedwyr. Both proofs in Abella and examples in Bedwyr are available online at https://github.com/kyagrd/NonBisim2DF.

Acknowledgment

Thanks to Sam Staton for providing an example that helped us in the completeness proof. The authors receive support from MOE Tier 2 grant MOE2014-T2-2-076. The third author receives support from NTU Start Up grant M4081190.020.

References

[1] D. Sangiorgi, "A theory of bisimulation for the π-calculus," *Acta Informatica*, vol. 33, no. 1, pp. 69–97, 1996.

[2] U. Montanari and V. Sassone, "Dynamic congruence vs. progressing bisimulation for ccs," *Fundamenta informaticae*, vol. 16, no. 2, 1992.

[3] R. Milner, J. Parrow, and D. Walker, "A calculus of mobile processes, Parts I and II," *Information and Computation*, vol. 100, no. 1, 1992.

[4] ——, "Modal logics for mobile processes," *Theoretical Computer Science*, vol. 114, no. 1, pp. 149–171, 1993.

[5] A. Tiu and D. Miller, "Proof search specifications of bisimulation and modal logics for the π-calculus," *ACM Transactions on Computational Logic*, vol. 11, no. 2, pp. 13:1–13:35, Jan. 2010.

[6] D. Baelde, K. Chaudhuri, A. Gacek, D. Miller, G. Nadathur, A. Tiu, and Y. Wang, "Abella: A system for reasoning about relational specifications," *J. Formalized Reasoning*, vol. 7, no. 2, pp. 1–89, 2014.

[7] D. Baelde, A. Gacek, D. Miller, G. Nadathur, and A. Tiu, "The Bedwyr system for model checking over syntactic expressions," ser. LNAI, F. Pfenning, Ed., no. 4603. Springer, 2007, pp. 391–397.

[8] G. Nadathur and D. Miller, "An Overview of λProlog," in *Fifth International Logic Programming Conference*. MIT Press, August 1988.

[9] A. Gacek, D. Miller, and G. Nadathur, "A two-level logic approach to reasoning about computations," *J. Autom. Reason.*, vol. 49, no. 2, 2012.

[10] M. Gazda and W. Fokkink, "Modal logic and the approximation induction principle," *Mathematical Structures in Computer Science*, vol. 22, no. 2, pp. 175–201, 2012.

[11] J. Parrow, J. Borgström, L. Eriksson, R. Gutkovas, and T. Weber, "Modal logics for nominal transition systems," in *CONCUR 2015*, ser. LIPIcs, vol. 42, 2015, pp. 198–211.

[12] U. Frendrup, H. Hüttel, and J. N. Jensen, "Modal logics for cryptographic processes," in *Proceedings of EXPRESS 2002*, ser. ENTCS, vol. 68, 2002.

[13] M. Abadi and A. D. Gordon, "A calculus for cryptographic protocols: The spi calculus," in *Fourth ACM Conference on Computer and Communications Security*. ACM Press, 1997, pp. 36–47.

[14] A. Tiu, "A trace based bisimulation for the spi calculus: An extended abstract," in *APLAS*, ser. LNCS, vol. 4807. Springer, 2007.

[15] A. K. Simpson, "The proof theory and semantics of intuitionistic modal logic," Ph.D. dissertation, University of Edinburgh, UK, 1994.

[16] D. Baelde, A. Gacek, D. Miller, G. Nadathur, and A. Tiu, *The Bedwyr System for Model Checking over Syntactic Expressions*. Springer Berlin Heidelberg, 2007, pp. 391–397.